

ND Prefix

Robustness Improvements

draft-vv-6man-nd-prefix-robustness

Paolo Volpato paolo.volpato@huawei.com

Eduard Vasilenko vasilenko.eduard@huawei.com


Background

- This draft looks at those cases that may lead to network prefixes invalidity
- Specific focus is on the multi-homing, multi-prefix scenario, even if the other cases are analyzed for the sake of completeness
- The target is to develop a root cause analysis and propose a solution.

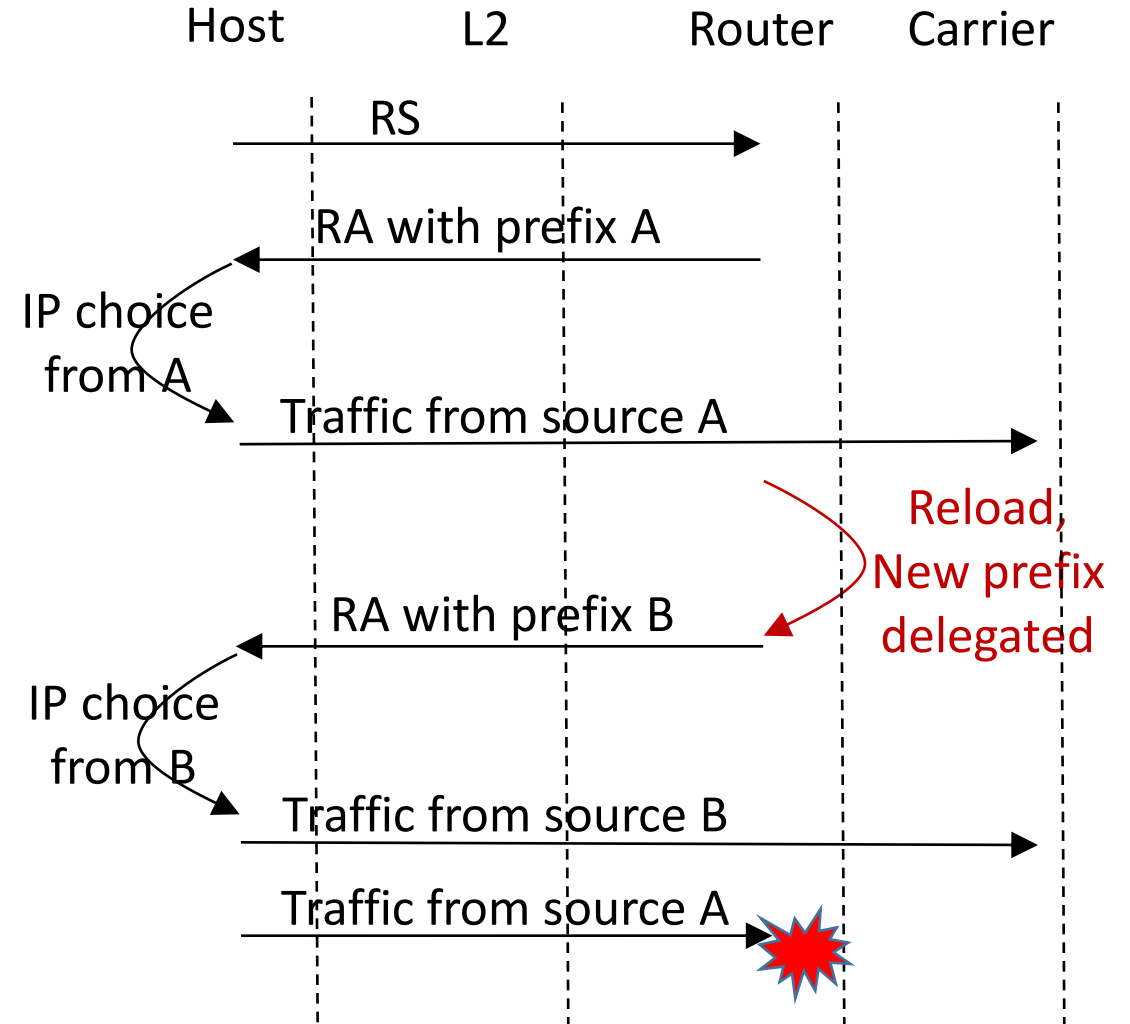
The Problem Statement – Router Reload (in an environment with Provider-Aggregatable addresses)

1A. Non-graceful reload

1B. Graceful reload, but prefix is not deprecated

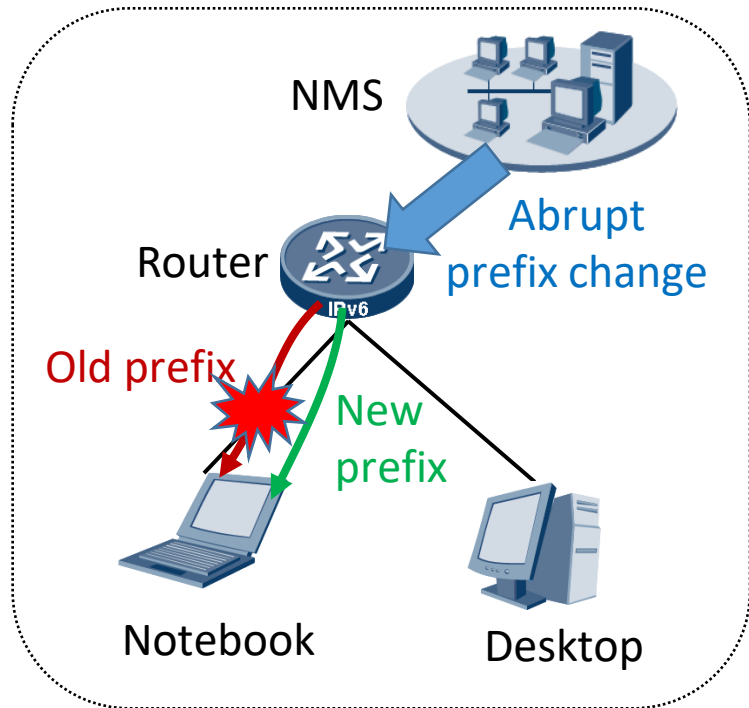


- Carrier could give new prefix
- Hosts do not see router's power recycle – could continue to use old prefix



The Problem Statement - Non-graceful Configuration Change (in any environment)


2A. Abrupt prefix change on the router



- Hosts do not receive prefix deprecation - old prefix could be used
- New prefix announced

New Prefix adoption would happen automatically in the SLAAC environment, DHCP adoption depends on host. Both cases makes it possible to use Old Prefix.

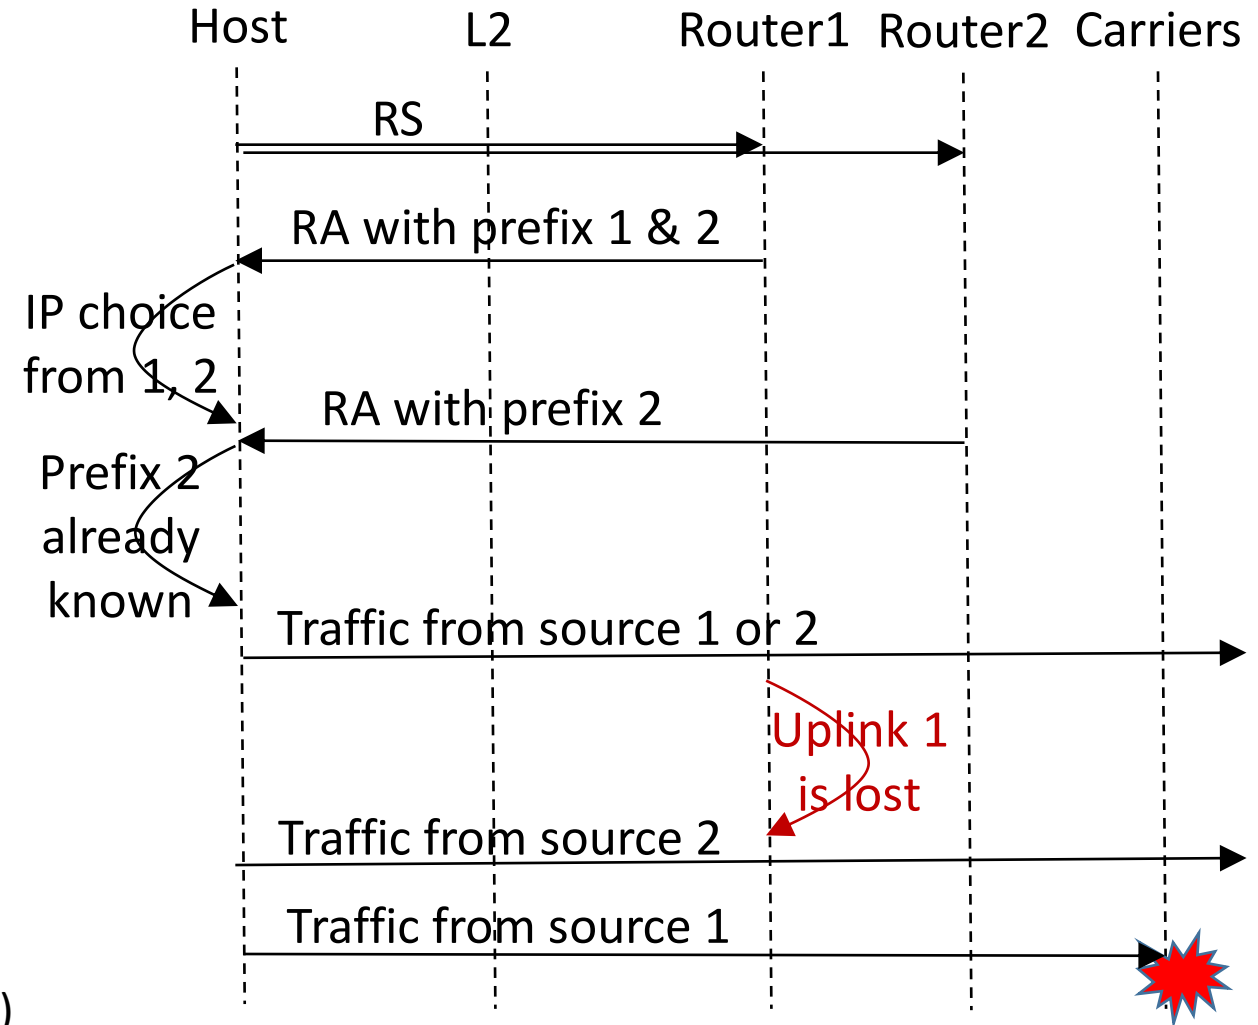
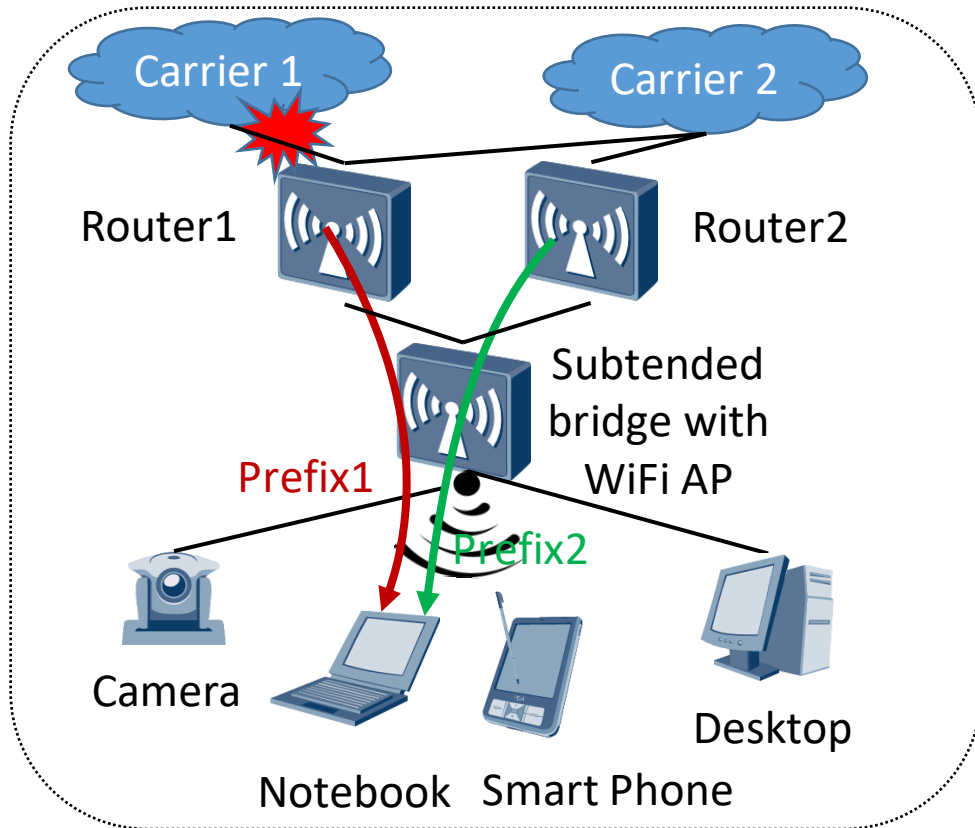
2B. VLAN change on the switch



- Hosts do not receive prefix deprecation – old prefix could be used
- New prefix (from different VLAN) announced

The Problem Statement – Multi-homing Multi-prefix environment

3A. Site connectivity if uplink is lost, but prefix is not deprecated automatically

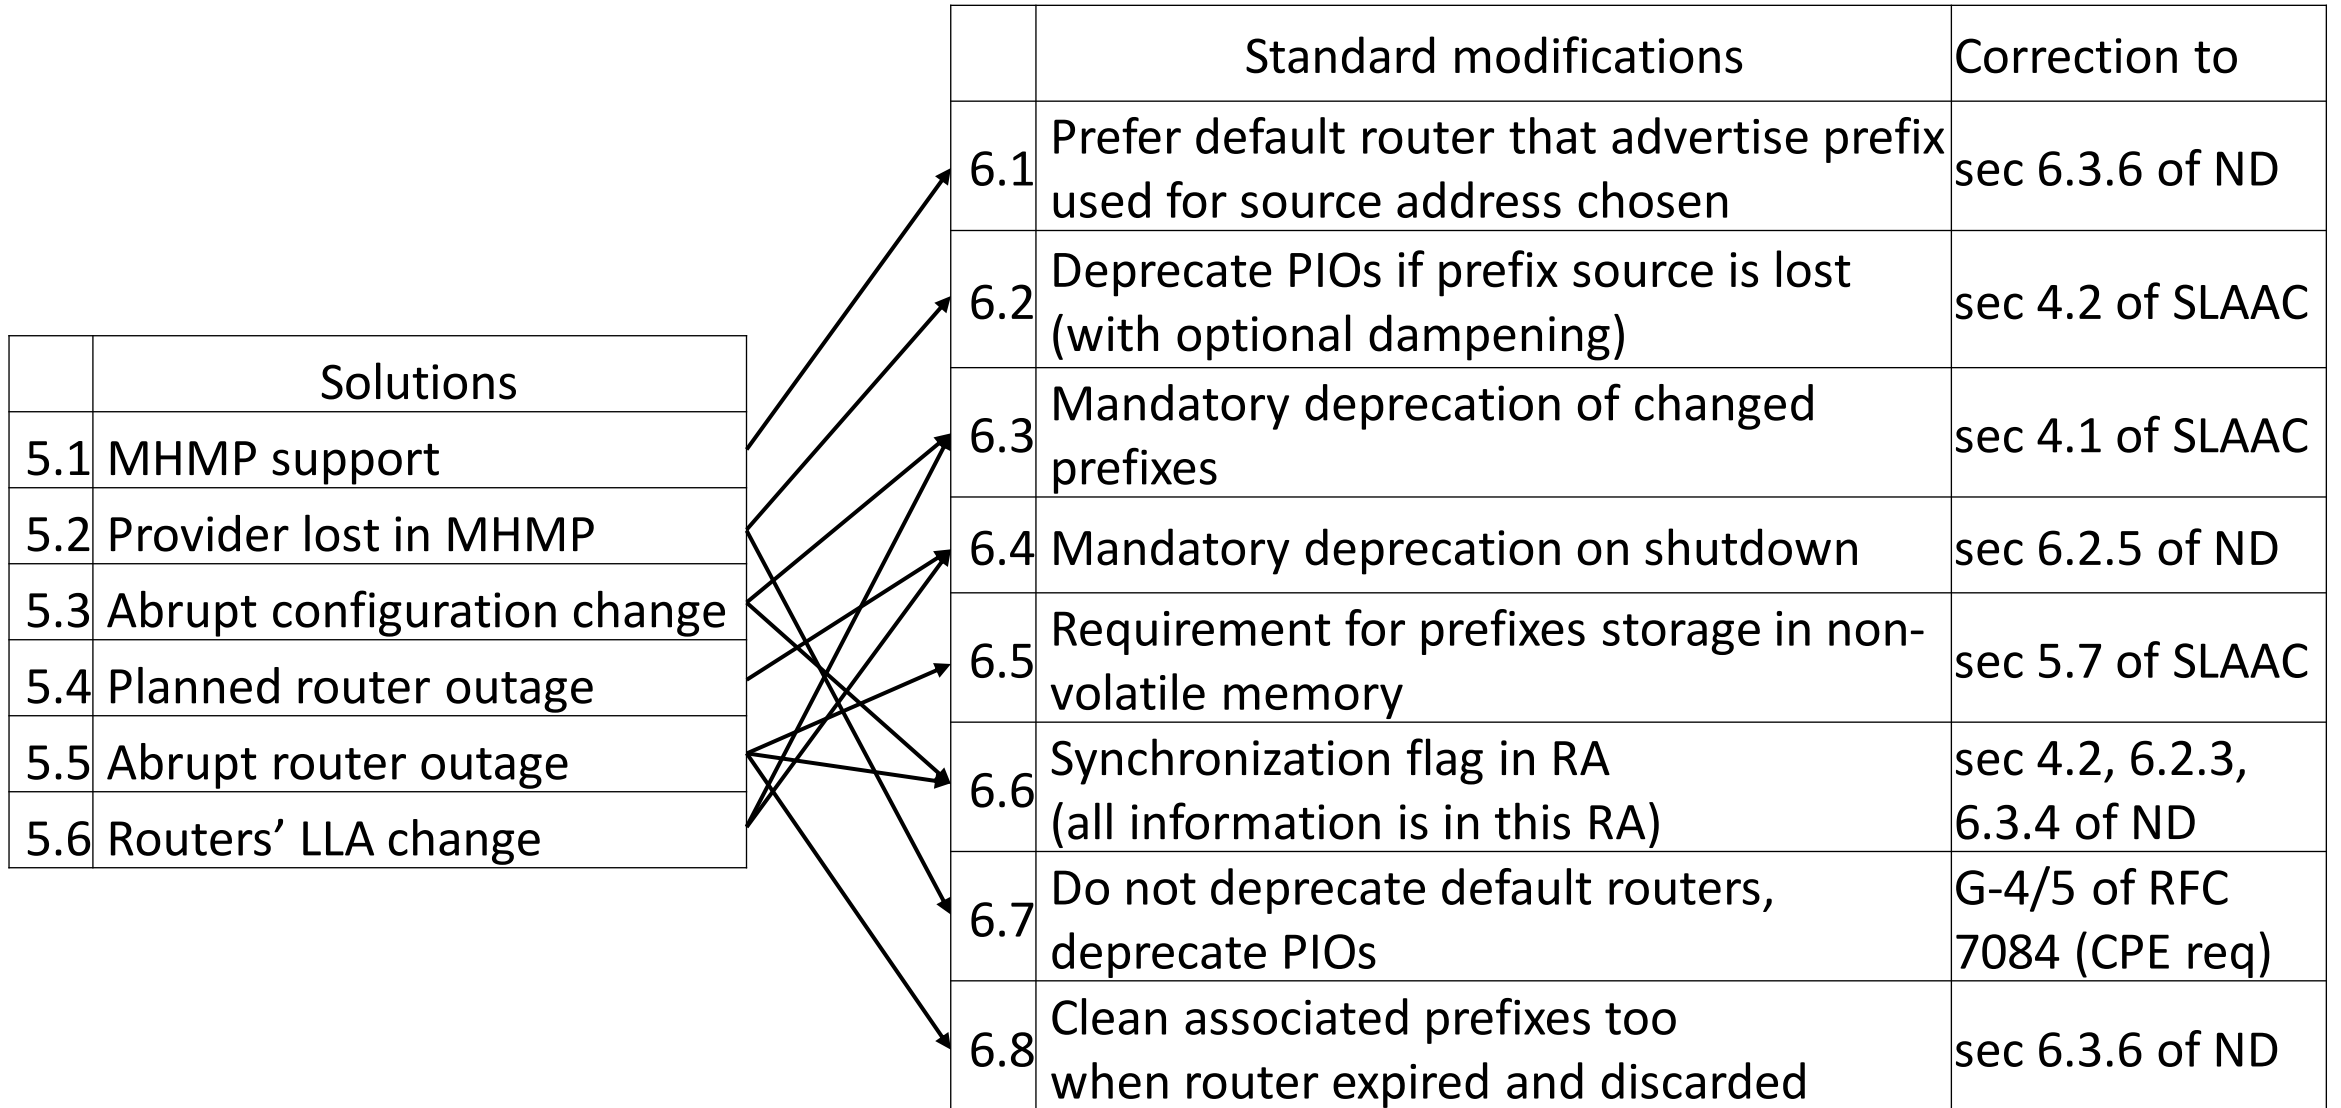


- Prefix1 and Default Router status was not deprecated because Router1 has connectivity to Carrier2 (RFC 7084)
- Hosts could use Prefix1, traffic would go to Carrier2

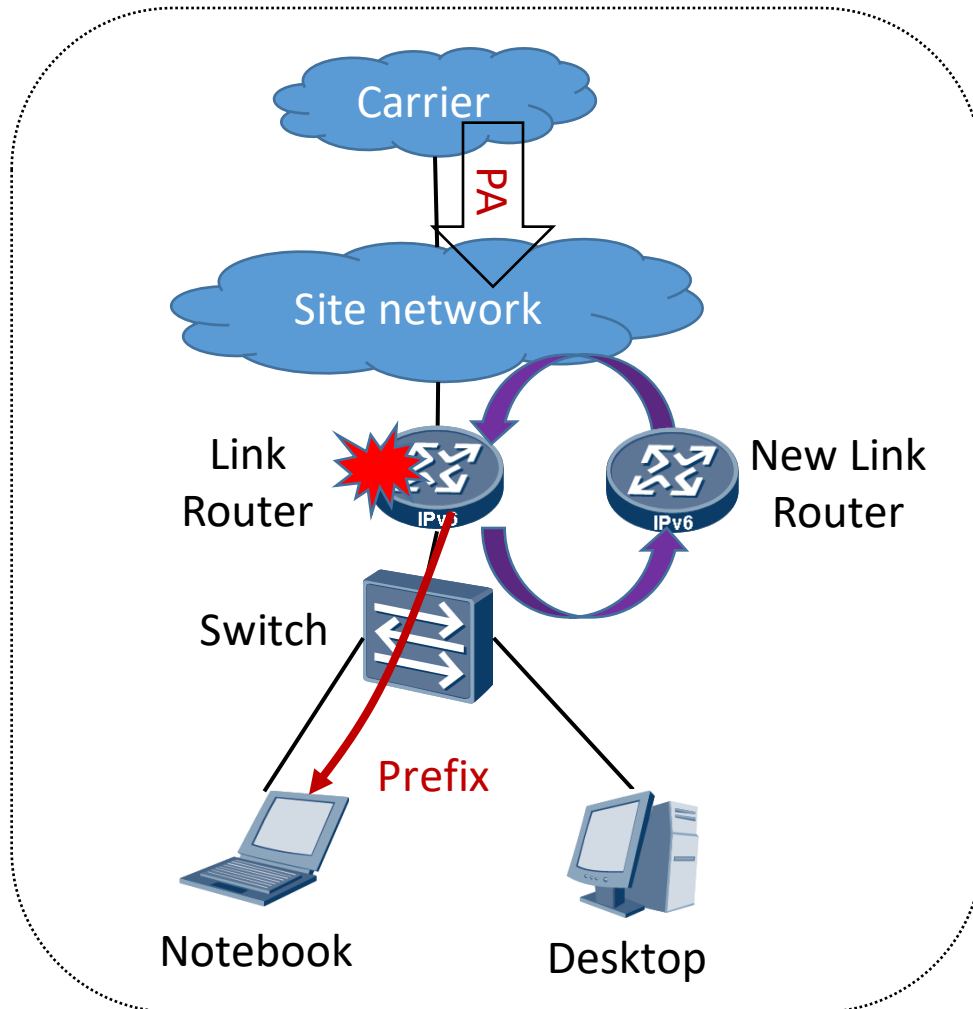
Result: traffic dropped by RPF check (for days in SLAAC or hours in DHCP)

Not discussed in RFC 8978

Solutions dependency on standard corrections



Possible improvement for the abrupt router replacement (spotted by Olorunloba Olopade from Virgin Media)



The problem statement:

Power could be abruptly disconnected

Then the router could be physically replaced.

Any announcement from the different router is not the reason for the host to clean the stale information of the previous router.

Potential solutions:

1. draft-olopade-6man-slaac-signaling: after any address promoted to preferred, host's clarification multicast RS request (to all active routers) sourced from GUA; router could analyze GUA and deprecate it immediately if the prefix is not accurate.

Drawbacks:

- many multicast checks after any temporary address promotion
- assumption that routers have always the same list of prefixes
- new address may not be asked in DHCP environment, the mechanism would not be activated

2. After any new router announcement, host's clarification multicast RS request; not responding router should be deprecated.

Drawbacks: 13s delay ($\text{MAX_RTR_SOLICITATIONS} \times$

$\text{RTR_SOLICITATION_INTERVAL} + \text{MAX_RTR_SOLICITATION_DELAY}$)

Next Steps

- We'd like to receive comments from the community
- Any review (even co-authoring) is welcome

- The next version of the draft will also redefine the introductory sections

- Thank you


Backup slides

ND topologies

Architecture I:


L2 and L3 devices are merged

- Sharing the fate for power, reboot



Architecture II: L2 and L3 devices are separate

- Multi-homing multi-prefix subnet is possible
- It is out of the scope of 1st hop how routers learn prefixes from upstream – any combination of prefixes could be announced from any router with possibility to change at any time

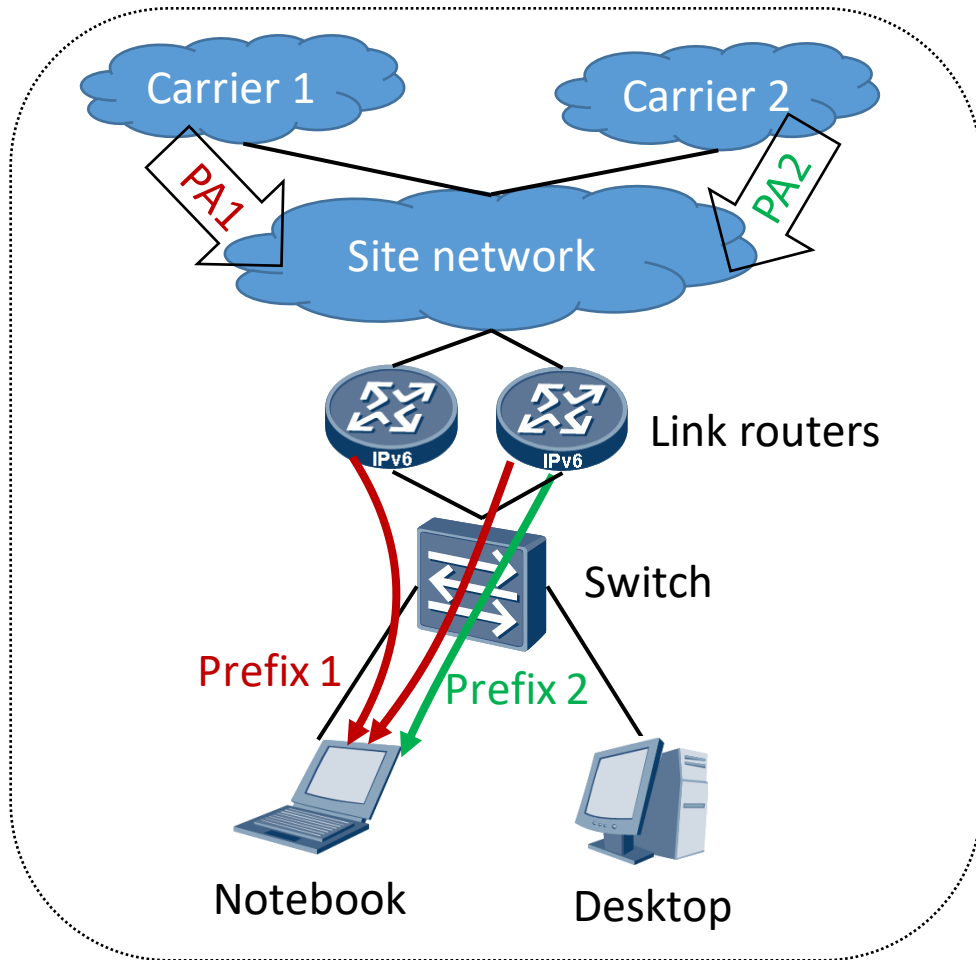


Protection Scenarios (assumed to be the full list)

1. Proper prefix usage for Multi-Homing Multi-Prefix environment.
Hosts should be capable of choosing in a coordinated way
(1) a source address (from proper PA prefix) and (2) a next hop:
 - a) In a normal situation: all providers and prefixes are available
 - b) In a faulty situation: one provider is not reachable, but some hosts and links on the routed path to this provider may still be reachable
 - c) In the case when the administrator abruptly replaces prefix
2. Prefix deprecation for the case of router outage that:
 - a) Planned for this interface
(reboot, shutdown, or ceasing to be a router)
 - b) Abrupt (power outage, software, or hardware bug)
3. Prefix deprecation for the case of link layer address change of the router
4. Prefix deprecation after abrupt router physical change

Bigger scope compare to RFC 8978 – MHMP is included

Multi-Homing Multi-Prefix: Challenge in Stable Environment



Reminders:


1. Any combination of PA prefixes could be announced from any router on the link. Router not announcing particular prefix may not have connectivity to respective Carrier.
2. It is assumed (by RFC 6724) in the majority of implementations that next hop is chosen first by default.

The problem for hosts: To coordinate the choice of next hop with source address to get connectivity.

Solutions:

1. The case “equal value prefixes” is resolved by RFC 6724: Default SASA (Source Address Selection Algorithm).
- 2a. The case “non-equal value prefixes” is partially resolved by: SASA (RFC 6724) with policy downloaded by DHCP (RFC 7078) + (should be coordinated) Route Preferences (RFC 4191)
- 2b. It is possible to resolve the case of “non-equal value prefixes” by only SASA (RFC 6724) with policy downloaded by DHCP (RFC 7078), If routers would be excluded from next hop selection for the cases when respective routers do not announce already chosen source IP address. **Modification to ND (RFC 4861) is proposed.**

Multi-Homing Multi-Prefix: Path to Delegated Provider is Lost



The same assumptions are as on the previous slide.

But the additional event is assumed: the path to 1 prefix is lost.

The problem for the router: Inform hosts.

The problem for hosts (the same): To coordinate the choice of next hop with source address to get connectivity.

Proposed solutions:

1. Withdraw prefixes, not default router status.

Modifications are proposed to:


- (1) CPE requirements (RFC 7084),
- (2) SLAAC (RFC 4862)

2. Join CPE requirements (RFC 7084) and HomeNet Architecture (RFC 7368) to use ULA for stability, security, and long outage protection. ULA could be the only address space available for a fully disconnected site.
3. Use a dampening mechanism to suppress oscillating prefix delegation uplinks.


Modification to SLAAC (RFC 4862) is proposed.

Non-graceful configuration change

On the router



On the L2 (switch) link

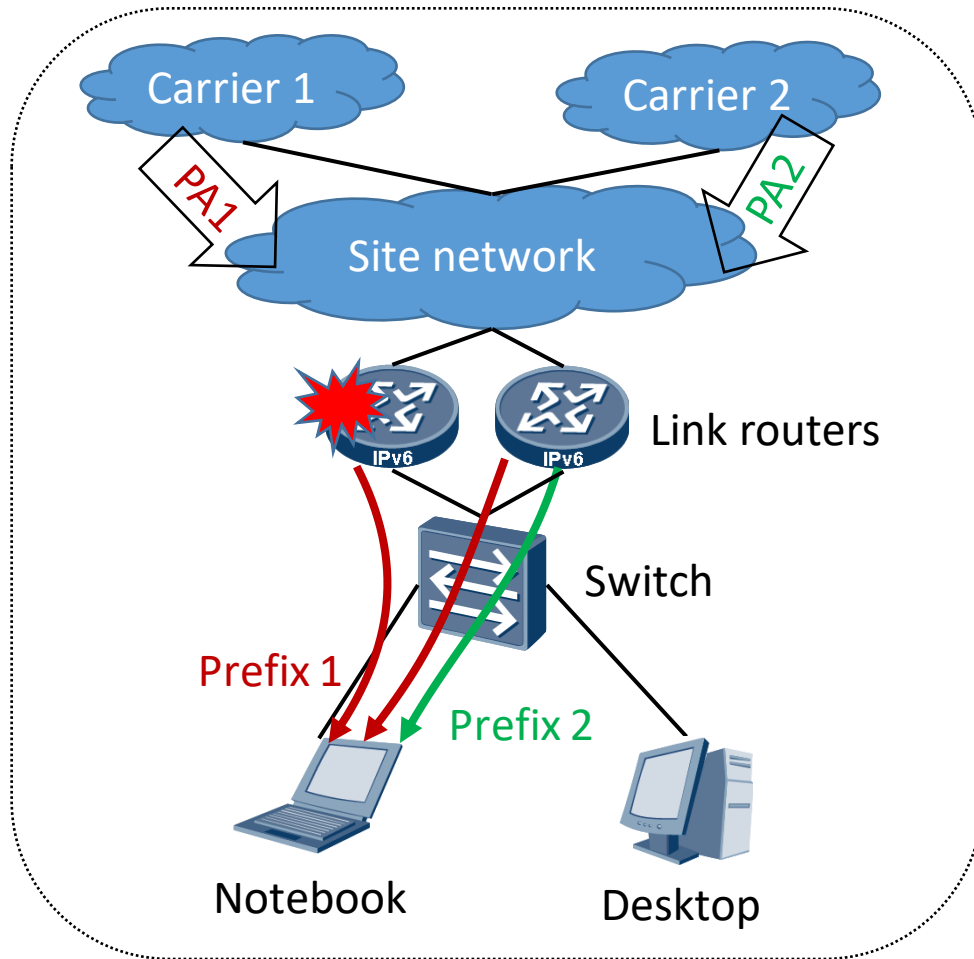


The problem statement: the host would use the wrong IP if moved to a different link (subnet) without any warning.

Proposed solutions:

1. Short term: additional check to make sure that prefix would be deprecated.
Modification to SLAAC (RFC 4862) is proposed
2. Long term: states synchronization between hosts and router.
Modification to ND (RFC 4861) is proposed

Router Outage



The problem statement:

router could use a different prefix for links after reload.

Proposed solutions for planned outage:

additional check to make sure that prefix would be deprecated.

Modification to ND (RFC 4861) is proposed

Proposed solutions for abrupt outage:

1. Short term: prefix storage in non-volatile memory

Modifications are proposed to:

(1) CPE requirements (RFC 7084, (2) SLAAC (RFC 4862)

2. Long term: states synchronization between hosts and router.

Modification to ND (RFC 4861) is proposed